SANS

Analyst Program

Whitepaper

A DNS Security Architecture as SecOps Force Multiplier

Written by **John Pescatore**

February 2023

Infoblox

Introduction

Domain Name System (DNS) services are like the central nervous system of digital business, enabling rapid and reliable communications. As such, DNS is both a highly attractive target for attackers and a key weapon in the continual struggle by security operations to quickly detect and repel those attackers. Providing better protection for DNS services reduces the likelihood of business disruption. Taking advantage of the vast amount of timely and accurate threat data available from DNS services means attacks that do get through will be detected more quickly and the focus can be on damage avoidance vs. incident response.

What Is a DNS Security Architecture?

DNS is essentially the internet's nervous system, providing and managing connections from the inside to the outside world. At the simplest level, DNS services manage and map human-understandable terms such as domain names to network IP addresses (i.e., 12.34.56.78). Without such services, there would be no connections between companies and their customers. Attackers know this and, as a result, DNS services are both a target of attackers and a resource they use to execute their attacks, exfiltrate data, and evade detection.

As shown in Figure 1, DNS is typically implemented across several servers or services used by the enterprise. These include the following:

- **DNS client—**DNS lookup queries originate within the enterprise from software (sometimes called a resolver) that takes user input (from PCs and other devices) or device input (from appliances or IoT devices) and then queries locally cached information obtained from a previous query or by connecting to a remote DNS server.
- **Local recursive DNS server—**The local DNS server takes the client query and checks with other DNS servers to find the correct IP address to return to the client. To speed response times, local DNS servers cache every response and can return results immediately if the cached value is deemed to be "fresh."
- **Internet forwarder—**If a DNS query cannot be resolved locally, an internet forwarder will query external trustable DNS servers to find the correct IP address for the client.
- **Authoritative DNS server—**Authoritative DNS servers store the most current and accurate DNS records for a given set of domains and provide the definitive mapping from user DNS lookup queries to IP addresses.

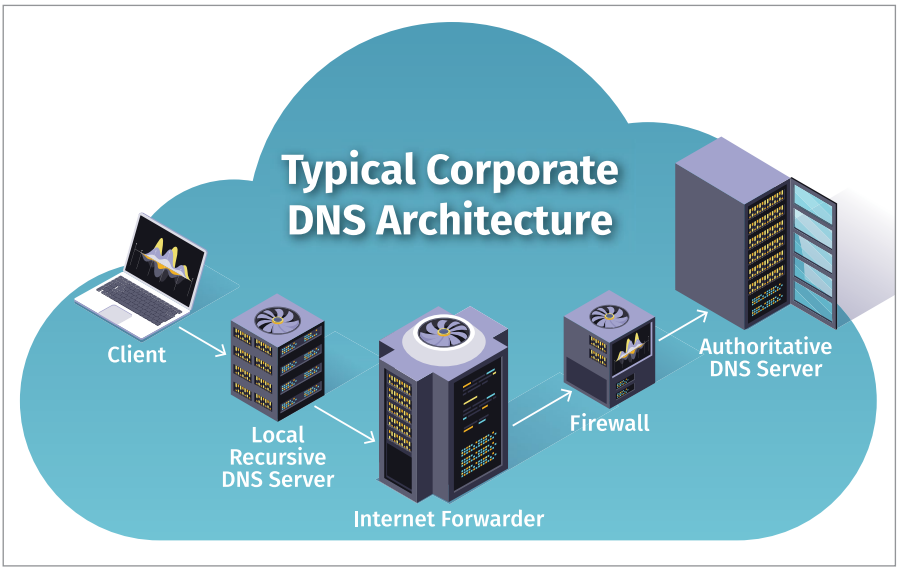

Figure 1. Typical Corporate DNS Architecture

Attackers are quite familiar with the components of DNS services for business use. If they can disrupt DNS queries and responses, they can cause a denial-of-service outage. If attackers can compromise the DNS process, they can cause malicious IP addresses to be returned to the unsuspecting clients and facilitate ransomware and other attacks.

A DNS security architecture is needed both to assure the integrity and availability of DNS services and to take advantage of DNS data and visibility to reduce the risks of a wide variety of attacks.

Protecting DNS Services

At a minimum, DNS queries that need to be answered by an external authoritative DNS server should pass through a firewall to minimize attack aperture. This firewall can be part of the enterprise perimeter firewall or can be a dedicated firewall protecting DNS services from external and internal attacks or misuse.

For DNS to function, however, the firewall obviously has to allow DNS traffic on port 53. Attackers realize that and for several years have used DNS tunneling techniques that take advantage of this to evade detection by intrusion detection/prevention systems as well as data loss prevention systems that are often deployed on other ports and protocols. To be effective, the firewall used here should support application-level filtering.[1]

Because DNS services are critical for business connectivity, they are an attractive target for distributed denial of service (DDoS) attacks.[2] There are two broad classes of DDoS attacks:

- **Brute force/flooding—**Attackers overload internet-connected servers to cause them to crash or to completely consume internet bandwidth.
- **Resource starvation—**Attackers connect to servers and initiate actions, such as site-wide searches or new-user registration processes, that completely consume available CPU or memory and block user access.

There are some denial-of-service techniques that aim to disrupt DNS services, such as NXDOMAIN attacks. NXDOMAIN is the response that a DNS server returns if there is no valid IP address for the domain name supplied. An NXDOMAIN DoS attack is a form of resource starvation and uses many DNS clients submitting queries for nonexistent or invalid domains. This attack causes the DNS server to consume its memory and CPU resources doing DNS recursion and sending NXDOMAIN responses, preventing legitimate DNS responses from being delivered to actual users.

NXDOMAIN attacks can be mitigated by detecting and blackholing or rate-limiting malicious or misbehaving domains, servers, or clients or by temporarily reducing timeout values for recursive lookups.

---

[1] A firewall is just one part of an overall strategy to protect DNS services. A good primer about protecting your DNS services can be found here: https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-81-2.pdf

[2] NIST SP 800-189 provides high-level guidance on DDoS detection and mitigation approaches. See https://csrc.nist.gov/publications/detail/sp/800-189/final

Using DNS to Increase Overall Security

Threat intelligence information provides details on the tactics, techniques, and procedures used by attackers. Often, threat intelligence includes lists of known malicious executables, domain names, and IP addresses. As this information comes out, enterprises can update block lists, but this is a reactive approach; attackers quickly and continuously switch domain names to avoid such simple filtering.

Taking advantage of the deeper visibility provided by the DNS system can support more proactive security practices, leading to improvements in three critical security metrics:

- Time to detect
- Time to respond
- Time to restore

The goal is to move as much as possible from incident response to damage minimization. All too often, businesses are not even aware of attacks until customers or law enforcement bring a compromise to their attention. More proactive precautions are necessary to prevent attacks and more quickly detect the start of unpreventable attacks. See Figure 2.

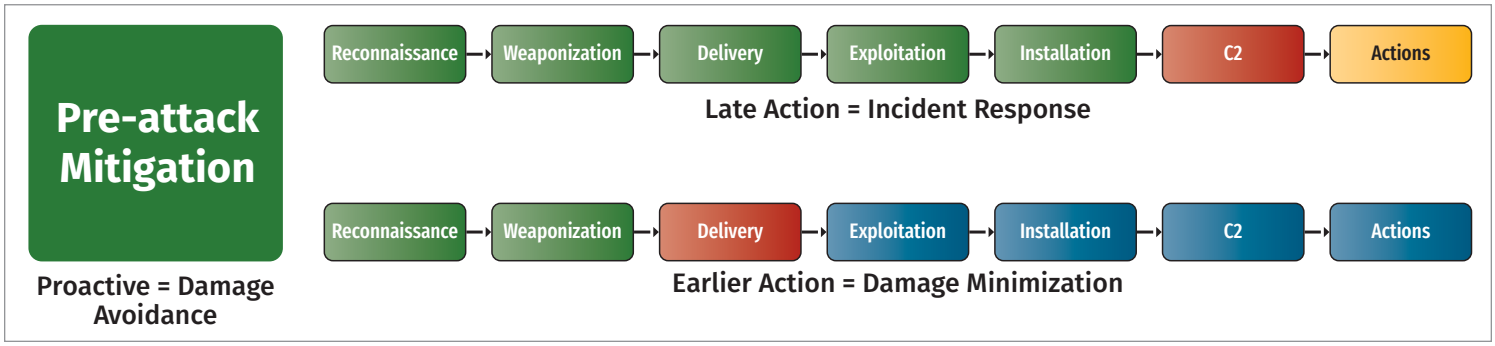

Figure 2. Pre-attack Mitigation

Because DNS requests are generally first seen quite early in the attack chain, DNS-based threat intelligence can be effective in shaping proactive precautions, reducing time to detect, and reducing the load on follow-on security controls. This also applies to appliances, IoT, ICS, and other devices without client-side visibility because they have to participate in DNS.

To be effective, all threat intelligence information must meet three key criteria that lead to measurable improvements in two or more of those three metrics:

- **Accuracy—**Minimal false positives
- **Timeliness—**Available before attacks succeed
- **Adaptability—**More than just reactive signatures

Threat intelligence consisting of standard signature lists of known malicious IP domains and IP addresses can score highly in accuracy with low false-positive rates. Because that approach is reactive and narrow, however, it fails miserably in the timeliness and adaptability categories. Behavior-based analysis and intelligence addresses adaptability, but if done manually, it usually fails the timeliness test. If done badly, it also fails the accuracy test. The use of machine learning techniques in threat intelligence has been touted to address both of these issues, but verifying those claims is difficult.

DNS response policy zones (RPZ) are an open standard that has been widely adopted to support the interchange of DNS firewall policy and configuration information. DNS threat intelligence providers track the domains that show evidence of serving up malicious content or participating in other ways through questionable activities. This "reputation" data can be distributed through the DNS architecture and used by DNS firewalls to implement trigger/response actions that can block or mitigate many forms of attack.[3]

Machine learning (ML) is a subset of AI that is usually defined by academics as "the field of study that gives computers the ability to learn without explicitly being programmed."[4] In reality, a lot of software needs to be written to be able to make use of ML at all, and a lot of good data needs to be collected and time spent training algorithms to make ML effective. Machine learning is not effective at simply ingesting random data and finding security-relevant information. DNS-based threat intelligence also has the advantage of providing high-fidelity data that can increase the effectiveness and accuracy of ML models.

Unsupervised learning is good for grouping similar events (clustering) that can be reviewed by skilled security analysts to determine whether a cluster defines a malicious behavior or is irrelevant to security decisions. Most use of ML in cybersecurity has been in "supervised learning" (see Figure 3), where known good input and output data are fed to ML tools, which then develop models that can classify unknown data as likely good or likely not good.

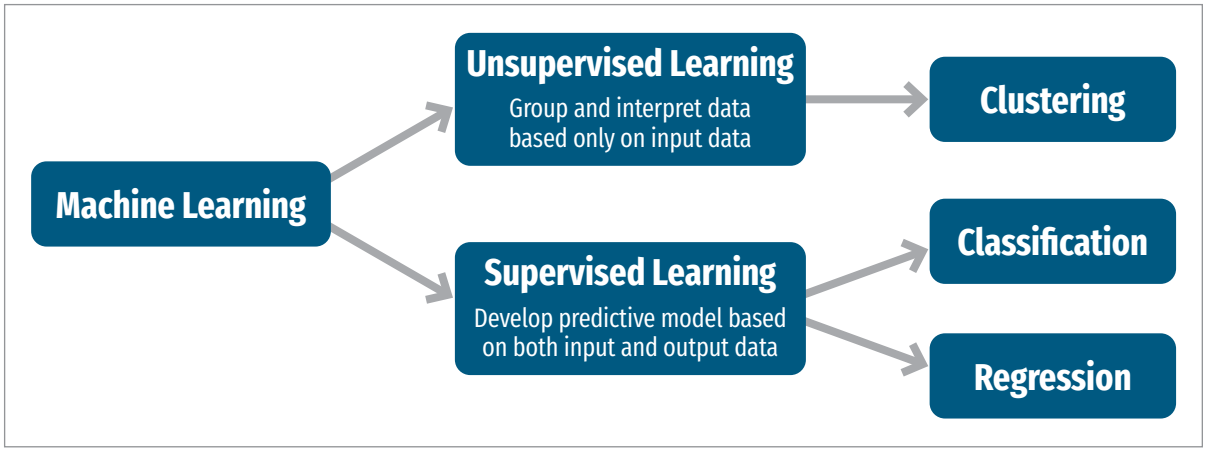

Figure 3. Supervised and Unsupervised Machine Learning Use

[3] Detailed information on DNS RPZ is available at https://datatracker.ietf.org/doc/html/draft-vixie-dns-rpz-04 and www.infoblox.com/dns-security-resource-center/dns-security-solutions/dns-security-solutions-response-policy-zones-rpz/

[4] "Machine Learning, Explained," https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained

In general, most successful use of machine learning in real world cybersecurity has been supervised learning augmented by reinforcement learning, where a skilled security analyst provides feedback on the "goodness" of the models' results. Security analysts who understand how DNS services are implemented and deployed can greatly magnify the improvements when reinforcement learning is used. This use of ML essentially enables the ML model to mimic the expert vs. autonomously "learn."

Because any vendor can claim to be using "machine learning," all claimed use of ML should be judged on some measurable improvement in time to detect, respond, and/or restore.

Example Use Case of Machine Learning

As far back as 2007, attackers have used DNS tunneling to evade detection.[5] DNS tunneling is a technique that takes advantage of data fields in DNS communications to embed attacker command and control traffic. This technique enables transversal of corporate firewalls that generally must allow DNS traffic to flow. Although the achievable transfer rate is usually in the low hundreds of kilobits range, DNS tunneling can also be used to exfiltrate data, evading data loss prevention (DLP) and other content-inspection approaches that only look at HTTP, Telnet, and FTP traffic.

Because there are many different approaches to carrying out DNS tunneling, simple signature-based detection approaches are not effective. By training machine learning models on known good DNS traffic, however, ML can be used as a high accuracy, low false positive automated detection technique for DNS tunneling.

Simple mitigation approaches, such as blacklisting or sinkholing the IP addresses of the DNS clients, can backfire when legitimate PCs have been compromised by attackers to launch the NXDOMAIN attack. But sinkholing generates significant additional traffic and overhead, and attackers using hardcoded DNS server and IP addresses can easily evade sinkholes.

There are several advantages to blocking on DNS:

- Blocking can occur early in the attack chain before damage occurs.
- Blocking facilitates higher fidelity identification of the infected endpoint.
- Appliances, ICS, IoT, and other devices usually can't host endpoint security software, but they have to participate in DNS.

The DNS system supports more flexible ways of reducing the likelihood of a client connecting to a malicious site, improving all three metrics. One example is RPZ.

The DNS system can also support trust decisions by providing more information than just the IP address. Data such as how recent the registration is, what registrar was used, the nameserver where the domain is registered, and what external associations exist can be used to implement more aggressive and flexible security policies.

[5] "Detecting DNS Tunneling," www.sans.org/white-papers/34152/

Best Practices for an Effective DNS Security Architecture

An effective DNS security architecture starts by ensuring the performance, availability, and integrity of DNS services by protecting the DNS host platform (server operating system, file system administrative apps and tools), the DNS software (name server, resolver), and the DNS data (zone file, configuration file). Essential security hygiene such as the CIS Critical Security Controls Implementation Group 1 (configuration management, patching, privilege management, etc.) are required for IT and security operations functions. Industry-secure configuration standards (such as Center for Internet Security [CIS] benchmarks and Department of Defense Security Technical Implementation Guides [DoD STIGs]) should be applied and audited for operating systems, databases, and DNS software. Widely accepted, broader frameworks such as the NIST Cybersecurity Framework and the MITRE ATT&CK® knowledge base provide the higher-level requirements and justifications for ensuring that a quality DNS security architecture is in use.

These best practices should also include using the DNS architecture for broader overall security benefits. When DNS services are secure and reliable, they can provide the key data for threat intelligence and attack detection/prevention/response capabilities discussed in the previous sections. DNS threat intelligence can provide early, accurate, and actionable information that supports thwarting attacks without causing inadvertent self-inflicted disruption. To gain these benefits, DNS should be an integral part of security operations, which often requires the SecOp staff (including both defenders and incident responders/investigators) to work closely and be cross-trained with IT or network ops groups that may have functional responsibility for DNS services. Where possible, common tools should be used across both groups.

Summary

A secure DNS architecture benefits business by ensuring the reliable and trustable DNS services needed for digital business. It also minimizes the risk of attacks compromising those same services, disrupting business and attacking customers. A well-designed and managed DNS architecture, combined with DNS threat intelligence, can reduce the "noise" produced by false-positive indications, reducing the load on security operations staff.

Sponsor

SANS would like to thank this paper's sponsor:

Infoblox®